

REMARKS

Claims 1-70 are pending in the application. Claims 7, 36 and 46 have been amended and claims 49-70 have been canceled.

Applicants believe that this response addresses the Examiner's rejection and that any changes do not introduce new matter into the specification, limit the scope of the claims or result in any prosecution history estoppel.

Claim Discussion – 35 USC §112

The Examiner rejected claims 1-70 under 35 USC §112, first paragraph, as based on a disclosure that is not enabling. In particular, the Examiner asserted that, for example, claim 1 "sorting memory transactions into at least one queue" is not enabling. Applicants respectfully disagree with the Examiner. In particular, the application would enable one of ordinary skill in the art to functionally replicate the disclosed element of the invention without likely substantial effort and/or experimentation. For example, the application at page 10, lines 3-10 provides:

In step 503, read transactions are sorted into queues based upon selected attributes. For example, read transactions are sorted into queues based upon the read address and address mapping used by memory. Read transactions in each queue are sorted by their arrival time. Other attributes including, but not limited to, CPU priority and demand versus prefetch transaction type may be used in sorting the transactions as well. Reads in separate queues will not have page conflicts with each other. A locally oldest transaction in each queue and a globally oldest transaction among queues may be identified.

The Examiner also rejected claims 7, 36 and 46 under 35 USC §112, second paragraph. In particular, the Examiner noted that "such as" renders the claim indefinite. Applicants respectfully note that the claims have been amended to delete the term "such as."

Claim Discussion – 35 USC S.103

The Examiner rejected claims 1-2, 4-5, 30-31, 33-34, 42, 44, 37-38 and 47-48 under 35 U.S.C. §103(a) as unpatentable over Foster (U.S. Patent No. 5,948,081) in view of Jenne, et al. (U.S. Application Publication No. 2002/0065981). Applicants respectfully note that neither Foster or Jenne, alone or in combination, teach or suggests the entire at least one queue is considered while making scheduling decisions as claimed or similarly claimed.

As noted in various parts of the present application, including on page 3, line 26 to page 4, line 14, embodiments of the invention provide for scheduling read and write transactions to memory out of order. In particular, memory transactions are sorted into queues and within each queue, memory transactions are sorted by arrival time. Memory transactions are scheduled out of order in accordance with read and write scheduling algorithms to optimize latency. Out of order selection is done at the time of launching a dram command (just in time scheduling). By making out of order selection at the time of dram command launch and not earlier, read and write queues are allowed to fill up without delaying the dram command. Filled queues result in more entries to choose from for an out of order scheduling decision, and thus better scheduling decisions can be made. *Moreover, the entire queue is considered while making scheduling decisions. Embodiments of the out of order scheduling policy are flexible enough to accommodate any page closing policy with no assumption being made about pages being in a particular “page miss/page empty/page hit” state.*

Foster fails to teach or suggest the entire at least one queue is considered while making scheduling decisions as claimed or similarly claimed. In particular, Figure 3 and column 3 lines 45 to column 4, line 7 (cited by the Examiner) are silent in this regard:

Broadly speaking, the present invention contemplates a computer system. The computer system comprises a memory requester, such as a CPU, graphics accelerator, or peripheral device (i.e., a device peripheral to a motherboard on which the CPU resides). The computer system further includes a memory, such as DRAM or SDRAM. An interface unit is coupled between the memory requester and the memory, and includes a first request queue and a second request queue. A queue controller is coupled between the first and second requests queues to determine, during operation, when a current memory dispatched from the memory requester to the second request queue is at the same memory address as a previous memory request stored within the first request queue. A memory controller is coupled to the queue controller for performing memory request to the memory bus. Read requests are serviced first unless (i) the write request queue requires flushing due to an address hit, or (ii) the write request queue is almost full. In the first instance (even though the write request queue is nearly empty), if there is a write request pending to the same address (i.e., a write hit) which a read request is to be serviced, then the write request must be de-queued or flushed first. In the latter instance, the write request queue must be de-queued in order before servicing the read request. When the write request queue is almost full, n memory requests are serviced from one request queue before toggling to the other, and then bursting n memory request from the other request queue. This provides a ping-pong flush technique of like n number of requests from each request queue.

Additionally, Jenne fails to teach or suggest making out of order selection at the time of DRAM command launch. The Examiner relies on Figure 8. Applicants respectfully note that neither figure 8 nor the accompanying specification directed to figure 8 noted below teaches or suggests making out of order selection at the time of DRAM command launch:

Referring still to FIG. 8, the memory controller 100 also preferably includes refresh logic 265, which according to normal convention, periodically issues a refresh command to particular memory addresses or sections of memory. According to the preferred embodiment, the refresh commands are stored in a refresh queue 260 prior to execution. The refresh commands are issued periodically to each section of memory to insure that data is not lost in DRAM.

[0060] The memory controller also preferably includes current calibration logic 275 that periodically issues a current calibration cycle to particular sections of memory to determine if the output current of the DRAM devices is within an acceptable range. According to the preferred embodiment, the current calibration cycles are stored in a queue 270 pending execution. Similarly, temperature calibration logic 255 also is provided in the memory controller 100 to periodically issue temperature calibration requests. These temperature calibration requests preferably are stored in a temperature calibration queue 250. Although each of logic devices 265, 275, 255 are depicted with separate queues in FIG. 8, they also may be implemented with queues that are integrated with the logic. As yet another alternative, these logic devices may be implemented without queues, by simply sending an appropriate request to the Rambus controller, which then determines the appropriate address for the refresh or calibration cycle.

[0061] According to the preferred embodiment of FIG. 8, arbitration logic 200 preferably couples to each of queues 210, 220, 230, 240, 250, 260 and 270 to determine which of these queues have pending transaction that are waiting to be executed. According to the preferred embodiment, the arbitration logic 200 selects or schedules a queue to proceed with executing a transaction. This monitoring and selection by the arbitration logic may be implemented in any of a variety of ways, as will be apparent to one skilled in the art. Thus, the queues may transmit a request signal to the arbitration logic, and the arbitration logic may then provide a grant signal back. Alternatively, the arbitration may snoop the output lines of the queues, and may provide a select signal to a multiplexer that passes the selected queue output for execution. Various other implementations are possible, and thus may be used in practicing the present invention. Whichever queue is selected by the arbitration logic 200, the highest priority transaction in that queue is then passed to the Rambus controller 105, which formats the request into an appropriate RDRAM protocol for transmitting to the targeted RDRAM location.

[0062] While certain aspects of the present invention have been described in relation to FIG. 8, it should be understood that the system shown in FIG. 8 is the preferred embodiment, not the only embodiment. Various designs may be used

to implement the principles of the present invention. Thus, any system that includes a high priority queue and at least one lower priority queue, and which runs calibration cycles, may benefit from the principles of the present invention. Thus, the specific cascading queue design of FIG. 8 is not a requirement of the present invention.

The Examiner rejected claims 1-2, 4-5, 30-31, 33-34, 42, 44, 37-38 and 47-48 under 35 U.S.C. §103(a) as unpatentable over Foster (U.S. Patent No. 5,948,081) and Jenne, et al. (U.S. Application Publication No. 2002/0065981) and further in view of Wulf, et al. (U.S. Patent No. 6,154,826). Applicants respectfully note the claims are patentable for the reasons noted above.

The Examiner rejected claims 1-2, 4-5, 30-31, 33-34, 42, 44, 37-38 and 47-48 under 35 U.S.C. §103(a) as unpatentable over Foster (U.S. Patent No. 5,948,081) and Jenne, et al. (U.S. Application Publication No. 2002/0065981) and further in view of Saracovsky, et al. (U.S. Patent No. 6,378,049). Applicants respectfully note the claims are patentable for the reasons noted above.

Applicants respectfully traverse the Examiner's rejection for the same reasons noted above. Additionally, there is also no motivation to combine these references.

CONCLUSION

In view of the foregoing, it is respectfully asserted that all of the claims pending in this patent application are in condition for allowance.

The required fee for a three month extension of time is enclosed. No additional fees are required for additional claims. Should it be determined that an additional fee is due under 37 CFR §§1.16 or 1.17, or any excess fee has been received, please charge that fee or credit the amount of overcharge to deposit account #02-2666.

If the Examiner has any questions, he is invited to contact the undersigned at (323) 654-8218. Reconsideration of this patent application and early allowance of all the claims is respectfully requested.


Respectfully submitted,
BLAKELY, SOKOLOFF, TAYLOR, & ZAFMAN LLP

Dated: May 15, 2006
By 
Farzad E. Amini, Reg. No. 42,261

12400 Wilshire Boulevard
Seventh Floor
Los Angeles, California 90025
(310) 207-3800

CERTIFICATE OF MAILING

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail with sufficient postage in an envelope addressed to: Mail Stop Amendment, Commissioner for Patents, Post Office Box 1450, Alexandria, Virginia 22313-1450 on May 15, 2006.


Margaux Rodriguez May 15, 2006